

Star Formation with Adaptive Mesh Refinement Radiation Hydrodynamics

Mark R. Krumholz¹

¹Dept. of Astronomy & Astrophysics, University of California, Santa Cruz, Interdisciplinary Sciences Building, Santa Cruz, CA 95064, USA
email: krumholz@ucolick.org

Abstract. I provide a pedagogic review of adaptive mesh refinement (AMR) radiation hydrodynamics (RHD) methods and codes used in simulations of star formation, at a level suitable for researchers who are not computational experts. I begin with a brief overview of the types of RHD processes that are most important to star formation, and then I formally introduce the equations of RHD and the approximations one uses to render them computationally tractable. I discuss strategies for solving these approximate equations on adaptive grids, with particular emphasis on identifying the main advantages and disadvantages of various approximations and numerical approaches. Finally, I conclude by discussing areas ripe for improvement.

Keywords. hydrodynamics, methods: numerical, radiative transfer, stars: formation

1. Introduction

While gravity is the dominant force in star formation, radiative processes are crucial as well. Most basically, radiation removes energy, allowing a collapsing cloud to maintain a nearly constant, low temperature as its density and gravitational binding energy rise by many orders of magnitude. Radiative cooling is what makes star formation a dynamic rather than a quasi-static process, a point understood quite early (see the review by Hayashi 1966). At densities below $\sim 10^{4-5}$ H atoms cm^{-3} , cooling occurs primarily via collisionally-excited atomic or molecular line emission, while at higher densities dust and gas become thermally well-coupled via collisions, and thermal radiation by dust grains is the dominant cooling process (e.g. Goldsmith 2001). At yet higher densities the coupled dust-gas fluid becomes optically thick to its own cooling radiation, causing the temperature to rise in a quasi-adiabatic fashion (Larson 1969).

While one-dimensional simulations of this collapse can include quite sophisticated treatments of radiative transfer between gas parcels (e.g. Masunaga et al. 1998), the computational difficulty of handling RHD in three dimensions led to a long period in which most 3D simulations of star formation simply dispensed with radiation entirely, choosing instead to represent the cooling processes either by a cooling function that removes energy at a rate dependent only on local gas properties, or by an even simpler equation of state that prescribes the gas temperature as a function of density. However, this approach has a major flaw: it does not allow energy exchange between parcels of gas, or between stars and diffuse gas. The last few years have seen a number of studies pointing out that this energy exchange is crucial to the dynamics of star formation, and that it can be dealt with only by RHD simulations.

The most important radiative transfer processes can be broken into three rough categories: thermal feedback, in which collapsing gas and stars heat the gas and thereby change its pressure; force feedback, in which radiation exerts forces on the gas that alter its motion; and chemical feedback, in which radiation changes the chemical state of the

gas (e.g. by ionizing it), and this chemical change somehow affects the dynamics. Each of these processes has been the object of intense study in the last few years.

With regard to thermal feedback, Larson (2005) pointed out that how gas fragments, and thus the stellar mass distribution that results from fragmentation, is extremely sensitive to how the temperature varies with density. Krumholz (2006) and Krumholz & McKee (2008) showed that one consequence of this sensitivity is that radiation produced by the first few stars to form in a given region can heat the gas around them, reducing the ability of that gas to fragment, favoring monolithic collapse to massive stars and suppressing formation of low mass objects. Subsequent numerical simulations (Krumholz et al. 2007a, 2010; Bate 2009; Offner et al. 2009; Urban et al. 2010) using a variety of methods have confirmed this effect.

Force feedback becomes important in the context of massive stars, which produce such high luminosities that the radiative force they exert on dusty interstellar gas can exceed their gravitational force. A number of authors have argued based on one-dimensional models that this process sets an upper limit on stellar masses (Kahn 1974; Yorke & Kruegel 1977; Wolfire & Cassinelli 1987). More recent work in two (Nakano 1989; Nakano et al. 1995; Jijina & Adams 1996; Yorke & Bodenheimer 1999; Yorke & Sonnhalter 2002; also see the contribution by Kuiper et al. in these proceedings) and three (Krumholz et al. 2007a, 2009) dimensions instead suggests that force feedback does not limit stellar masses. Deciding the question requires radiation-hydrodynamic simulations.

The most important type of chemical feedback is photoionization, which raises the gas from ~ 10 K to $\sim 10^4$ K. This causes a number of important dynamic effects, including ejecting mass from star-forming clouds (Dale et al. 2005; Peters et al. 2010), altering the magnetic field structure (Krumholz et al. 2007c), and driving turbulent motions (Gritschneder et al. 2009). Radiation also drives other chemical processes, but these are usually not important for dynamics, and thus may be handled by post-processing simulations. In this review I limit myself to radiative effects that are dynamically important and must therefore be simulated in tandem with the hydrodynamic evolution of the system.

2. The Equations of Radiation Hydrodynamics

The equations of radiation hydrodynamics (RHD) in conservation form read (Mihalas & Weibel-Mihalas 1999)

$$\frac{\partial}{\partial t} \begin{pmatrix} \rho \\ \rho \mathbf{v} \\ \rho e \end{pmatrix} + \nabla \cdot \begin{pmatrix} \rho \mathbf{v} \\ \rho \mathbf{v} : \mathbf{v} + P \\ (\rho e + P) \mathbf{v} \end{pmatrix} = \begin{pmatrix} 0 \\ \mathbf{G} \\ cG^0 \end{pmatrix}, \quad (2.1)$$

where ρ , \mathbf{v} , e , and P are the gas density, velocity, specific energy, and pressure. The source term on the right hand side represents the rate at which radiation transfers momentum and energy to the gas. For simplicity I have omitted terms describing gravity and magnetic fields, which would appear as additional sources on the right hand side.

The rate at which radiation transfers energy and momentum to the gas is described by the radiation four-force vector (G^0, \mathbf{G}) ,

$$c \begin{pmatrix} G^0 \\ \mathbf{G} \end{pmatrix} = \int d\nu \int d\Omega [\kappa_\nu(\mathbf{n}) I_\nu - \eta_\nu(\mathbf{n})] \begin{pmatrix} 1 \\ \mathbf{n} \end{pmatrix} \quad (2.2)$$

where \mathbf{n} is a unit vector, I_ν is the radiation intensity at frequency ν in direction \mathbf{n} , and κ_ν and η_ν are the extinction coefficient and emissivity of the gas as a function of frequency

ν and direction \mathbf{n} . Finally, the intensity is governed by the transfer equation,

$$\frac{1}{c} \frac{\partial}{\partial t} I_\nu + \mathbf{n} \cdot \nabla I_\nu = \eta_\nu(\mathbf{n}) - \kappa_\nu(\mathbf{n}) I_\nu. \quad (2.3)$$

For simplicity I have omitted terms describing scattering, which is generally not important for the radiation-hydrodynamics of star formation.

The equations of RHD are, like the ordinary equations of fluid dynamics, characterized by dimensionless numbers. The two most important ones for radiation-hydrodynamics are $\beta \sim v/c$, where v is the characteristic value of \mathbf{v} , and $\tau \sim L/\lambda_p$, where L is the size-scale of the system and λ_p is the photon mean free path. The first ratio β characterizes how relativistic the system is. In writing equation (2.1) in non-relativistic form, we have already assumed that $\beta \ll 1$. The second, τ , characterizes how optically thick it is. Systems with $\tau \ll 1$ are described as being in the streaming limit, and are characterized by weak matter-radiation interaction. Those with $\tau \gg 1$ are in the diffusion limit and have strong-matter radiation coupling. Both limits appears in star formation problems.

3. Approximations and Solution Methods

The full system of equations (2.1) and (2.3) is seven-dimensional, with quantities varying in time, space (3 dimensions), direction (2 dimensions), and frequency. Unfortunately, this makes full numerical solution prohibitively expensive even on modern supercomputers. We are therefore reduced to approximations. There are two broad classes of approximation in common use in star formation simulations: moment methods and characteristic methods. A third category, Monte Carlo methods, has been used extensively in post-processing radiative transfer calculations, but has not been used extensively in radiation-hydrodynamic simulations. I will not discuss it in detail.

3.1. Codes

In the discussion that follows I introduce the most common methods used in AMR RHD methods, and in Table 1 I summarize which method(s) are used in each code. The codes I include in the table are as follows: **Orion** is the oldest and probably best-tested AMR RHD code, but is not publicly available, and currently lacks a magnetohydrodynamic (MHD) capability (Klein 1999; Howell & Greenough 2003; Krumholz et al. 2007b). **Ramses** is an AMR MHD code that was recently upgraded with a radiation solver, although the latter is not (as of this writing) publicly available (Fromang et al. 2006; Commerçon et al. 2010). **Flash** is an open source AMR MHD code with extensive chemistry capabilities and several add-on modules that handle radiation in different ways (Fryxell et al. 2000; Rijkhorst et al. 2006; Peters et al. 2010). **Enzo** is an open source AMR HD code (an MHD version exists but is not public) with cosmology capabilities and two different radiation methods (Abel & Wandelt 2002; Norman et al. 2008; Reynolds et al. 2009). **Pluto** is an open source AMR MHD code with a newly-developed radiation solver (Mignone et al. 2007; Kuiper et al. 2010). There are also a number of fixed, nested grid codes in use, but I will not mention these by name.

3.2. Moment Methods

3.2.1. Basic Theory of Moment Methods

The basic idea of a moment method is to take moments of the transfer equation, in exact analogy to the Chapman-Enskog procedure used to derive the equations of fluid dynamics from the kinetic theory of gases. To take the zeroth moment we integrate the transfer equation over all directions; to obtain the first moment we multiply both sides

Moment methods			Characteristic methods		
Code	ν resolution	Frame	Code	ν resolution	Ray scheme
Orion	MG, 2T	Mixed	Orion	1F	HEALpix
Ramses	2T	Comoving	Flash	MG	Hybrid characteristics
Enzo	2T	Comoving	Enzo	1F	HEALpix
Pluto	1T	Comoving	Pluto	MG	Sphere

Table 1. Summary of Codes. See text for explanation of fields. All moment codes use first order closure (flux-limited diffusion). Enzo and Orion appear in both lists because they can operate in characteristic or moment mode. Pluto appears on both lists because it uses a hybrid characteristic-moment method (see text for details). Orion has both a MG and a 2T moment method.

by the unit vector \mathbf{n} and then integrate; for the second moment we multiply by the rank two tensor $\mathbf{n} : \mathbf{n}$ and integrate, and so forth. As in the analogous fluid case, the procedure yields an exact solution if carried out to infinitely many orders, but one instead makes an approximation by truncating the procedure after finitely many orders. Usually “finitely many” here means one or two, and the first two moments of equation (2.3) are

$$\frac{\partial}{\partial t} \begin{pmatrix} E \\ \mathbf{F}/c^2 \end{pmatrix} + \nabla \cdot \begin{pmatrix} \mathbf{F} \\ \mathbf{\Pi} \end{pmatrix} = - \begin{pmatrix} cG^0 \\ \mathbf{G} \end{pmatrix}, \quad (3.1)$$

where

$$E = \frac{1}{c} \int d\nu \int d\Omega I_\nu \quad (3.2)$$

$$\mathbf{F} = \int d\nu \int d\Omega I_\nu \mathbf{n} \quad (3.3)$$

$$\mathbf{\Pi} = \frac{1}{c} \int d\nu \int d\Omega I_\nu \mathbf{n} : \mathbf{n} \quad (3.4)$$

are the first three moments of the radiation intensity: the radiation energy density, radiation flux, and radiation pressure tensor. In these equations the direction-dependence in the transfer equation is removed and the dimensionality of the problem is lowered by two, but at the price of introducing the radiation pressure tensor, for which we do not have an equation (since we have not expanded the transfer equation to the next order) and must instead make an approximation.

Because moment methods necessarily involve smoothing the angular dependence of the radiation field, they are best suited to representing diffuse, smooth radiation fields. This makes them ideal for handling thermal and force feedback, which tend to be dominated by diffuse infrared light re-radiated by dust grains. They are less suited for chemical feedback, which is usually dominated by direct, highly beamed radiation from a handful of stellar sources. A further advantage of moment methods is that their computational cost is independent of the number of sources, and usually scales only as N or $N \log N$, where N is the number of cells. They also parallelize fairly easily.

3.2.2. Types of Moment Methods

Moment methods can be classified based on several design choices that are made in their construction, all of which involve some tradeoff between computational expense and accuracy. See Table 1 for a list of the design choices made in each of the popular codes.

Closure order. The first choice is whether to retain both of the first two moment equations, or to retain only the zeroth moment and adopt a closure approximation for the radiation flux; all the moment methods currently used in practical star formation

simulations make the latter choice, using a closure known as flux-limited diffusion (FLD). In a medium where the photon mean free path is small compared to the size of the system, Eddington showed that the radiation flux in the fluid rest frame approaches

$$\mathbf{F} = -\frac{c}{3\kappa_R}\nabla E, \quad (3.5)$$

where $\kappa_R = \int d\nu (\partial B_\nu / \partial T) / \int d\nu \kappa^{-1} (\partial B_\nu / \partial T)$ is the Rosseland mean opacity, $B_\nu(T)$ is the Planck function. This is known as the diffusion approximation. While it is quite good in optically thick media, it breaks down in optically thin regions because as $\kappa_R \rightarrow 0$ the signal speed approaches infinity rather than properly limiting to c . The FLD approximation is an attempt to fix this problem by instead setting the radiation flux to

$$\mathbf{F}_0 = -\frac{c\lambda}{\kappa_R}\nabla E_0, \quad (3.6)$$

where λ is a function that has the property that $\lambda \rightarrow 1/3$ in optically thick regions and $\lambda \rightarrow \kappa_R E_0 / |\nabla E_0|$ in optically thin regions, so that $|\mathbf{F}_0| \rightarrow cE_0$, its correct limiting value. Many choices for the function λ are possible; the most common one in astrophysical applications is the Levermore & Pomraning (1981) limiter, $\lambda(R) = R^{-1}(\coth R - R^{-1})$, where $R = |\nabla E_0| / (\kappa_R E_0)$. However, all of these limiters are of unknown accuracy in the intermediate optical depth regime, and all FLD methods suffer from the problem that they discard information about the directionality and momentum content of the radiation field. Higher order moment methods exist and can solve some of these problems (e.g. variable tensor Eddington factor methods, Hayes & Norman 2003, and the M1 closure method, González et al. 2007), but none have yet proven cheap and robust enough for use in practical AMR star formation simulations.

Frequency resolution. The second choice in setting up an RHD moment method is the level of resolution in frequency. Most accurate is the multigroup (MG) method, in which one discretizes equation (3.1) in frequency, solving one version of the equation for each frequency bin. All the bins are coupled via the matter temperature, which affects the extinction and emissivity and thus the radiation four-force vector (equation 2.2). Next most accurate is the 2T method, in which one takes the radiation spectrum to be a Planck function characterized by a single temperature T_{rad} (thus removing a dimension from the problem), but allows T_{rad} to be different than the gas temperature T_{gas} . Simplest of all is the 1T method, in which one assumes $T_{\text{rad}} = T_{\text{gas}}$, allowing a further simplification of the equations. As one might expect, this choice involves a tradeoff between accuracy and cost; the 1T method is cheapest but badly misestimates both the radiation spectrum and T_{gas} in the streaming regime. In comparison 2T methods still produce incorrect radiation spectra in the streaming regime, but make much smaller errors in the matter temperature. They are intermediate in cost. MG provides a good representation of both the spectrum and the matter temperature but is most expensive.

Choice of frame. The third design choice is whether to formulate the equations in the comoving frame or using mixed frames. Extinction κ and emissivity η are simple, isotropic functions in the frame comoving with the fluid, but in other frames they contain complex directional dependence as a result of relativistic boosting. One might think that this effect is small in non-relativistic flows, but ignoring it is equivalent to neglecting the work done by radiation on the gas (Mihalas & Klein 1982), which is obviously unacceptable if radiation forces are non-negligible. To avoid complex velocity dependence in η and κ , it is highly desirable to write them in the comoving frame. The simplest choice after that is to write the radiation quantities (E , \mathbf{F}) in the comoving frame as well (e.g. see Mihalas & Klein 1982), but this carries a significant cost. Since the comoving frame

is non-inertial in any system with non-constant fluid velocity, the comoving radiation energy is not a conserved quantity, and thus comoving-frame codes cannot be exactly conservative. Moreover, every time the resolution changes and the velocity field is refined, the reference frame change as well, introducing errors in conservation that are likely to accumulate with increasing refinement. The alternative is a mixed-frame formulation in which one writes the radiation quantities in the inertial lab frame. Deriving this formulation is tricky, since one must carefully account for the Lorentz transformations between frames, but the resulting equations are explicitly conservative, and can be discretized in a manner that conserves energy to machine precision (Krumholz et al. 2007b).

3.3. Characteristic Methods

In characteristic methods one solves the transfer equation (2.3) directly, but only along selected rays. Given this solution, one can compute the radiation four-force vector (equation 2.2) directly. This provides much greater accuracy in computing the radiation intensity along those rays, but at the cost of neglecting all other rays. For this reason it is best used for chemical feedback, which tends to be dominated by the rays coming from a small number of sources. It is not well-suited to handling diffuse radiation fields, and simulations based on this technique are generally no better than non-radiative simulations when it comes to handling effects like fragmentation being altered by a diffuse IR radiation field. The cost of these methods scales as the number of cells times the number of sources, but with a coefficient that is generally smaller than for moment methods. Thus these methods are cheaper than moment methods when the number of sources and rays is small, but become impractically expensive if the number of sources is even a small fraction of the number of cells. As with moment methods, there are several design choices to be made in setting up a characteristic method.

Frequency resolution. As in the moment method case, one can choose either to retain the frequency-dependence in the transfer equation by using multiple frequency groups (the MG method), or one can integrate over frequency. In this case one does not generally assign a temperature to the radiation; instead, since for chemical feedback one cares about photons only above a certain energy, the usual approximation is to adopt an average photon energy. This is the single-frequency (1F) approximation.

Ray-drawing scheme. There are also a number of schemes for drawing rays from the point source(s). The simplest is the spherical ray technique: one uses spherical coordinates, requires that the source be at the origin, and simply draws rays aligned with the computational grid. This parallelizes almost perfectly and is very simple to code, but is obviously unsuited to any problem with multiple stellar sources or where the locations at which stars form is not known *a priori*. The other two schemes in wide use are HEALPix-based adaptive trees (Abel & Wandelt 2002) and the hybrid characteristics (Rijkhorst et al. 2006), both of which operate on top of an underlying Cartesian grid. The HEALPix scheme is based on the Hierarchical Equal Area isoLatitude Pixelization of the sphere, which divides the sphere into a equal area pixels that can be subdivided indefinitely, allowing the angular resolution to adapt to match that of the underlying grid. The hybrid characteristics scheme works by combining a method of short characteristics within individual AMR grids with a method of long characteristics between grids. Neither suffer from the limitations of the spherical ray method.

3.4. Hybrid Methods

Hybrid methods attempt to combine the best features of characteristic and moment methods. Characteristics work well for the highly beamed radiation coming directly from a star or some other point source. However, once this radiation is absorbed, following its

re-emission with characteristics becomes unreasonably expensive, and a moment method is a far better choice. The underlying idea of a hybrid method, first used in multidimensional simulations by Murray et al. (1994) and recently implemented in the Pluto code (Kuiper et al. 2010), is to use characteristics for the “first absorption”, then switch to a moment method. One does this by performing a characteristic trace from the star or stars to determine the rate of energy and momentum input into each cell by direct radiation. One adds this as a source term on the right hand side of the moment equation (3.1), then solves as one would in a pure moment method. The computationally-cheap characteristic step can use MG, while the more expensive moment solve uses the 1T or 2T approximation. Since the first absorption is generally the part of the problem where the frequency dependence is most important, this method achieves some of the accuracy of a fully frequency-dependent calculation at significantly lower cost.

4. Future Directions

I close this review with a brief discussion of possible future directions for AMR RHD in the star formation context. One such likely direction comes from combining realistic treatments of thermal, force, and chemical feedback into a single simulation, rather than treating only one or two of them at a time as in current simulations. Since the first two effects are most easily handled by moment methods and the third by characteristic methods, hybrid techniques are the natural solution. To be competitive in handling thermal and force feedback, however, these methods will have to be linked to more advanced moment methods than has been attempted before. Nonetheless, a hybrid method combining frequency-dependent characteristic tracing with a 2T, mixed frame FLD method is a straightforward extension of current techniques, and seems a logical place to start.

Further in the future, second order moment methods are likely to become important. Developing them to the point where they can run reliably in a parallel environment on adaptive grids will be a major algorithmic undertaking, one that is likely to require the assistance of computer scientists and applied mathematicians in addition to astronomers. If successful, these methods will be able to handle both beamed and diffuse radiation fields within a single framework, and will remove many of the limitations of the FLD approximation. Monte Carlo methods are another candidate to replace the current dominance of FLD and characteristic methods. They are very computationally costly, since many photons are required to suppress Poisson noise, but they have the advantage of near perfect parallelization. Since the future of supercomputing seems to be heading toward the development of ever-larger numbers of processors, each of which is no faster than the processors of the previous generation, this may prove to be a decisive advantage.

Finally, a caveat: for all the limitations of our current RHD methods, it is not entirely obvious that our numerical techniques are the limiting factor in the accuracy of our simulations. The main agent for coupling radiation and gas at the high densities where stars form is dust. Our knowledge of dust grain properties, such as sublimation temperatures and grain size distributions in dense environments, is quite limited, and appears likely to advance less quickly than either computer power or algorithmic technique. It may be the case in the future that our knowledge of dust becomes the limiting factor in the accuracy of RHD simulations.

References

- Abel, T., & Wandelt, B. D. 2002, *MNRAS*, 330, L53
 Bate, M. R. 2009, *MNRAS*, 392, 1363

- Commerçon, B., Hennebelle, P., Audit, E., Chabrier, G., & Teyssier, R. 2010, *A&A*, 510, L3+
- Dale, J. E., Bonnell, I. A., Clarke, C. J., & Bate, M. R. 2005, *MNRAS*, 358, 291
- Fromang, S., Hennebelle, P., & Teyssier, R. 2006, *A&A*, 457, 371
- Fryxell, B., et al. 2000, *ApJS*, 131, 273
- Goldsmith, P. F. 2001, *ApJ*, 557, 736
- González, M., Audit, E., & Huynh, P. 2007, *A&A*, 464, 429
- Gritschneider, M., Naab, T., Walch, S., Burkert, A., & Heitsch, F. 2009, *ApJL*, 694, L26
- Hayashi, C. 1966, *ARA&A*, 4, 171
- Hayes, J. C., & Norman, M. L. 2003, *ApJS*, 147, 197
- Howell, L. H., & Greenough, J. A. 2003, *JCP*, 184, 53
- Jijina, J., & Adams, F. C. 1996, *ApJ*, 462, 874
- Kahn, F. D. 1974, *A&A*, 37, 149
- Klein, R. I. 1999, *J. Comp. App. Math.*, 109, 123
- Krumholz, M. R. 2006, *ApJL*, 641, L45
- Krumholz, M. R., Cunningham, A. J., Klein, R. I., & McKee, C. F. 2010, *ApJ*, 713, 1120
- Krumholz, M. R., Klein, R. I., & McKee, C. F. 2007a, *ApJ*, 656, 959
- Krumholz, M. R., Klein, R. I., McKee, C. F., & Bolstad, J. 2007b, *ApJ*, 667, 626
- Krumholz, M. R., Klein, R. I., McKee, C. F., Offner, S. S. R., & Cunningham, A. J. 2009, *Science*, 323, 754
- Krumholz, M. R., & McKee, C. F. 2008, *Nature*, 451, 1082
- Krumholz, M. R., Stone, J. M., & Gardiner, T. A. 2007c, *ApJ*, 671, 518
- Kuiper, R., Klahr, H., Dullemond, C., Kley, W., & Henning, T. 2010, *A&A*, 511, A81+
- Larson, R. B. 1969, *MNRAS*, 145, 271
- . 2005, *MNRAS*, 359, 211
- Levermore, C. D., & Pomraning, G. C. 1981, *ApJ*, 248, 321
- Masunaga, H., Miyama, S. M., & Inutsuka, S. 1998, *ApJ*, 495, 346
- Mignone, A., Bodo, G., Massaglia, S., Matsakos, T., Tesileanu, O., Zanni, C., & Ferrari, A. 2007, *ApJS*, 170, 228
- Mihalas, D., & Klein, R. I. 1982, *JCP*, 46, 97
- Mihalas, D., & Weibel-Mihalas, B. 1999, *Foundations of Radiation Hydrodynamics* (Mineola, New York: Dover)
- Murray, S. D., Castor, J. I., Klein, R. I., & McKee, C. F. 1994, *ApJ*, 435, 631
- Nakano, T. 1989, *ApJ*, 345, 464
- Nakano, T., Hasegawa, T., & Norman, C. 1995, *ApJ*, 450, 183
- Norman, M. L., Bryan, G. L., Harkness, R., Bordner, J., Reynolds, D., O'Shea, B., & Wagner, R. 2008, in *Petascale Computing: Algorithms and Applications*, ed. D. A. Bader, Computational Science Series (Chapman & Hall/CRC), 83–101, arXiv:0705.1556
- Offner, S. S. R., Klein, R. I., McKee, C. F., & Krumholz, M. R. 2009, *ApJ*, 703, 131
- Peters, T., Banerjee, R., Klessen, R. S., Mac Low, M., Galván-Madrid, R., & Keto, E. R. 2010, *ApJ*, 711, 1017
- Reynolds, D. R., Hayes, J. C., Paschos, P., & Norman, M. L. 2009, *Journal of Computational Physics*, 228, 6833
- Rijkhorst, E., Plewa, T., Dubey, A., & Mellema, G. 2006, *A&A*, 452, 907
- Urban, A., Martel, H., & Evans, N. J. 2010, *ApJ*, 710, 1343
- Wolfire, M. G., & Cassinelli, J. P. 1987, *ApJ*, 319, 850
- Yorke, H. W., & Bodenheimer, P. 1999, *ApJ*, 525, 330
- Yorke, H. W., & Kruegel, E. 1977, *A&A*, 54, 183
- Yorke, H. W., & Sonnhalter, C. 2002, *ApJ*, 569, 846

Acknowledgements

I thank Richard Klein, Chris McKee, and Loius Howell for helpful discussions. I acknowledge support from: an Alfred P. Sloan Fellowship; NASA through ATFP grant NNX09AK31G and through the Spitzer Theoretical Research Program, through a contract issued by the JPL; and the NSF through grant AST-0807739.